

B1
cancel. FIG. 26 is a block diagram illustrating three SAN domains, each managed by a SAN manager and having heterogeneous disk arrays.

Insert the following paragraphs after the paragraph ending at page 21, line number 9, as follows:

The following devices can make the data delivery possible to the applications:

Storage Devices: Disk, Tape

Storage Controllers: RAID, Disk & Tape Controllers

Access Controllers: Server, Switch

Access Managers: Server, SAN Managers

Data End Points: Servers, Clients

B2
omit
10068352-0529001
2006250-25389001
The Storage devices are typically dumb magnetic platters that store the data. The storage controllers have intelligence and can be either host bus adapter or an external controller closer to the disk/tape. The access controllers today are the servers, which define the criteria for the access to data. The switch can perform this action in the future. The data access manager or the data manager is a client or a server today and can be the SAN managers or switches in the future. The data end points are applications that reside in the clients and the servers today.

FIG. 17 represents the typical storage stack in a controller/server or a client. Referring to FIGs. 17, 18 and 19, the disk and the tape drivers 1701 give the access to the physical data. The data can be stored on a RAID array 1801, 1901 or 1902 in case of which, the RAID layer 1702 will provide the access to the data. Multiple RAID arrays 1801, 1901 and 1902 or JBODs 1802, 1903 and 1904 can be clubbed together to form a stripe or a mirror 1803, 1905 and 1906. Many vendors like Auspex, Veritas, DEC, IBM support this model.

The volume manager today assumes that all the members of the volume are locally attached. However, with the increasing reliability and speeds of the LAN and WAN, these disks can be LAN or WAN attached.

FIG. 20 shows 2 servers 2001 and 2002 that are connected on a LAN/WAN 2005. Each server has a RAID array 2003 and 2004 attached to it. The RAID arrays on both ends are sliced equally into 2 pieces, A & B. Say the pieces are named A-1, B-1 & A-2 and B-2. A mirror can be defined on server-X that has A-1 and A-2 as its members and mirror on server-Y has B-1 and B-2

as its members. So when a write happens from server-X on A-1, the mirrored member A-2 also gets the data to be written, over the LAN/WAN. Today, many proprietary solutions exist to make this kind of topology happen. But each one of these solutions is different from one another and do not interoperate.

Also with the advent of SANs and increasing adaptation of the SAN technology, these disks in FIG. 20 can be SAN attached.

There is nothing that operates at the volume layer today, that can provide abstraction as to where the disks are physically attached. Also with the increasing need for interoperable solutions, this abstraction should support all the available standards of LAN/WAN/SAN. This abstraction can be provided using a new protocol that understands the semantics of LAN, WAN and SAN. Auspex Systems Inc. initiated the definition of this new protocol and named it VVMP- Virtual Volume Management Protocol. The protocol should work in the frameworks shown in FIG. 21.

As shown in FIG. 21, VVMP 2101 can be both in-band and out-of-band.

VVMP 2101 must take care of the horizontal storage methodologies and also on the vertical access methodologies, as shown in FIG. 22.

VVMP 2101 must take care of the following devices:

Devices that broadcast the storage parameters

Devices that support downloading of the SAN characteristics

Devices that are dumb (JBODs and Tapes)

Devices that support other specific protocols (HiPPI....)

VVMP 2101 must work with both:

FiberChannel

LAN/WAN Switches

VVMP 2101 should support the following data path/session management functions for servers and clients (applications).

Requests for storage with specific characteristics of type, capacity, bandwidth and QOS;
request for a change in these parameters at a later time

Mechanism for applications to register for storage allocation (on availability)

Requests for storage with exclusive or shared access

B2
Cm'
10068352-052902

Requests for additional storage or give away excess storage

Request for automatic backup and mirroring

Requests to include or exclude special hardware in the data path

Request to show statistics for a particular data session

Request to use RAID parameters like stripe size, cache line size, block size

Security Management by creating storage domains

Essentially VVMP 2101 is a Client-Server Protocol for Volume Management and setting up the Data Path - Data Path Management. FIG. 23 illustrates an example for the data path management.

In FIG. 23, an application running on a server 2301 is using the video data on a RAID volume and the data is being compressed on write and uncompress on read using a compression hardware that is on the SAN. The SAN management 2302 has been instructed to set up the data path and the participating devices (compression hardware 2303, server 2301 and the RAID controller 2304) are bound for the length of the data session.

Figure 24 is another instance of SAN/WAN interoperability using VVMP. Server A 2401 and Server B 2402 use a volume X 2403. Server A 2401 has the volume connected over a local SAN and Server B 2402 accesses the Volume X 2403 over a WAN/SAN Combination 2404.

FIG. 25 shows a generic computing subsystem that supports both SAN and NAS. The disks can be attached on the SAN or connected on a LAN/WAN using a router 2501. VVMP, though defined to be a data path management protocol, can encapsulate other protocols that can be used to setup the data path or manage it. E. g.

VVMP	NHRP for SAN
VVMP	DHCP for SAN

NHRP for SAN can be a protocol that helps the application or the SAN manager discover the storage attached on other SANs. DHCP for SAN can be a protocol that can be used by the

edge devices to download SAN specific characteristics for the device. FIG. 26 shows interconnected SANs.

FIG. 26 shows 3 SAN domains, each managed by a SAN manager 2601-2603 and having heterogeneous disk arrays 2604-2606. Each SAN has servers 2607-2609 and clients using the volumes that are created on the disk arrays 2604-2606. The following is an example of the VVMP protocol frame format.

OPCODE		Source SAN ID	
Source SAN ID	Destination SAN ID		
Client ID			
Client ID			
Server ID			
Server ID			
Transaction ID			
Sub OPCODE	Length of Data in 32bit words		Checksum

Sample OPCODES

0x0001 - Query Request

0x8001 - Query Reply

0x0002 - Allocate Request

0x8002 - Allocate Reply

0x0003 - Initialize Request

0x8003 - Initialize Reply

0x0004 - Free Request

0x8004 - Free Reply

Sample Sub OPCODES

For Query

0x01 - Query of SAN

0x02 - Broadcast for all devices

0x03 - Broadcast for local devices

0xFF - Broadcast of all devices on all SANs

For Allocate

0x01 - Local SAN only

0x02 - Allocate with facility for expansion

0xFF - Anywhere and all the above facilities

Sample frame format for Allocate

Storage Type		Access Type	Session ID
Capacity in KB			
Bandwidth in KB			
QOS			
Flags	Timeout in msec		

Storage Types

0x00 - JBOD

0x01 - RAID0

0x02 - RAID2

0x03 - RAID3

0x04 - RAID4

0x05 - RAID5

0xFF - Any type

Access Types

0x01 - SCSI

0x02 - FC

0x03 - HiPPI

0x04 - ATM

0xFF - Any type

Each data session will have a session ID, that is assigned by the VVMP server, which typically is the SAN manager or a switch. The session ID is valid for the time of the session, which can be torn down by the client or the server or on a reboot of either of them.

*B2
cncld.*
10068352
25589001

In FIG. 26, each volume manager (intelligent if it is a RAID/Tape controller or the SAN manager 2601-2603 itself in the case of JBOD) will broadcast its attributes periodically or in case of a query. The SAN Manager 2601-2603 will register it. In FIG. 25, when server-A 2607 requests a RAID5 volume of 100GB, the SAN-1 Manager 2601 will allocate the storage depending on the parameters requested by the server. When all the storage in SAN-1 is exhausted, the SAN-1 Manager 2601 can now go to other SANs to request the storage depending on the sub-opcode.

For the out-of-band management, VVMP will use a reserved port on which the server process will listen for requests.

28-36
Please delete Appendix A, pages ~~22-30~~, in the specification.

In the drawings:

N.E. Please add sheets 17-22 attached.